

REMARKS

Reconsideration and allowance are respectfully requested.

The Examiner raises several drawing objections. Regarding Figure 2A, the MEM1 block is already labeled 230, and not as element 23 as the Examiner suggests. A copy of Figure 2A as filed is attached with this response. Regarding the objection Figure 1, the specification at page 14, line 29, has been amended to include the underlined text: "Via paths 85 and 90, load instructions may be used to load data values into the register bank from the data memory 87, and store instructions may be used to store data values into the data memory 87 from the register bank 35." This amendment overcomes the Examiner's objection.

The Examiner also objects to Figure 2. But based on the content of the objection, it appears that the Examiner may have intended to refer to Figure 9 rather than Figure 2. As explained on page 28, lines 27 to 29, the instruction queue is a modified version of the standard queue, which incorporates an instruction decoder 205. Figure 9 shows how the instruction queue may be implemented. A standard queue is described earlier on page 26 with reference to Figure 6. A person skilled in the art would understand that those parts of Figure 9 which correspond with equivalent parts in Figure 6 would operate in the same way unless otherwise specified. But to remove any doubt, a final sentence has been added at page 29, line 8 to state explicitly that the buffer flags 940 and 950 and the multiplexers 960, 970 operate in the same manner as described earlier with reference to the buffer flags 640, 650 and multiplexers 660, 670 of Figure 6.

Withdrawal of the objections to the drawings is requested.

The Examiner objects to the length of the abstract. Amendments have been made so that the word count is less than 150 words. Withdrawal of this objection is requested.

All of the claims 1 to 45 stand rejected as being anticipated by US6,240,508 to Brown.

This rejection is respectfully traversed.

To establish that a claim is anticipated, the Examiner must point out where each and every limitation in the claim is found in a single prior art reference. *Scripps Clinic & Research Found. v. Genentec, Inc.*, 927 F.2d 1565 (Fed. Cir. 1991). Every limitation contained in the claims must be present in the reference, and if even one limitation is missing from the reference, then it does not anticipate the claim. *Kloster Speedsteel AB v. Crucible, Inc.*, 793 F.2d 1565 (Fed. Cir. 1986). Brown fails to satisfy this rigorous standard.

The present application is directed to data processing apparatus that has both a main processor and a coprocessor, where both the main processor and the coprocessor have a plurality of pipeline stages. Coprocessor instructions in a sequence of instructions to be executed by the data processing apparatus are routed to the coprocessor for execution. As explained in the background, each coprocessor instruction may be routed through both the main processor pipeline and the coprocessor pipeline, but this means that the main processor pipeline and the coprocessor pipeline need to be synchronized. Such synchronization may be achieved by passing signals with fixed timing from one pipeline to the other. As the length of pipeline processors increases, it is more difficult to achieve synchronization between pipelines using such a tightly coupled synchronization approach.

The inventors conceived of a synchronization approach that uses at least one synchronizing queue to couple a predetermined pipeline stage in one of the pipelines with a partner pipeline stage in the other pipeline. The predetermined pipeline stage causes a token to be placed in the synchronizing queue when processing a coprocessor instruction. The partner pipeline stage then processes that coprocessor instruction upon receipt of the token from the

synchronizing queue, thereby synchronizing the first and second pipelines at that point. This inventive approach allows some "slack" between the two pipelines so that strict synchronization at all stages is not necessary. At the same time, the pipelines are correctly synchronized for crucial transfers of information.

Brown is concerned with preserving read and write ordering in a macro-pipelined processor of the type that decouples instruction decode and instruction execution so as to allow multiple macroinstructions to exist in the pipeline at various stages of processing. See column 3, line 50, to column 4, line 14. Figure 1 shows a macro-pipelined processor 10 with an instruction unit 22 (referred to as the I-BOX), an execution unit 23 (referred to as the E-BOX), and a memory management unit 25 (referred to as the M-BOX). See column 7, lines 24 to 41. In the type of system described in Brown, maintaining read and write ordering is essential; otherwise, access to memory is not deterministic. Column 49, lines 51 to 57 explain that some memory requests may originate from the E-BOX 23 rather than from the I-BOX 22. As a result, these explicit memory requests must be synchronized with references from previous and subsequent instructions. This is achieved by providing a spec-queue 75 and a spec-queue sync counter within the M-BOX 25. The M-BOX 25 is described in more detail with reference to Figures 15 and 18. See also Brown's claim 1.

The Examiner appears to be equating the claimed main processor with Brown's CPU 10 and the claimed coprocessor with Brown's additional CPU 28 shown in Figure 1. But CPU 28 is not described in Brown as a coprocessor for the CPU 10. Column 35, lines 45 to 49, referred to by the Examiner, merely describe a memory ownership technique used to ensure one processor can write to a memory location before another processor can read from the same memory location. But a memory ownership technique is not what is recited in the independent claims.

But setting aside the issue of whether CPU 28 is a coprocessor, Brown's CPU 28 does not execute CPU 28 instructions appearing in the same sequence of instructions executed by the CPU 10.

The independent claims also recite that the coprocessor includes a second pipeline. Brown does not disclose that CPU 28 is pipelined. Claims 1 and 29 also recite that each coprocessor instruction is arranged to be routed through both the first pipeline (i.e., the pipeline of the main processor) and the second pipeline (i.e., the pipeline of the coprocessor). In contrast, the Examiner admits that Brown's instructions are issued to only one processor. Accordingly, the anticipation rejection is improper because specifically claimed features relating to the coprocessor in claims 1 and 29 are not disclosed in Brown.

Other claim features are missing from Brown. When addressing the claimed synchronizing queue, the Examiner refers to column 49, lines 58 to 60, which identifies the earlier-mentioned spec-queue 75 appearing within the M-BOX 25 of the CPU 10. The spec-queue 75 is provided entirely within the main processor 10 to synchronize when the E-BOX 23 can issue memory requests in addition to the I-BOX 22. Accordingly, Brown's spec-queue 75 cannot be equated with the claimed "at least one synchronizing queue," because the synchronizing queue "coupl[es] a predetermined pipeline stage in one of the pipelines with a partner pipeline stage in the other of the pipelines," and hence, couples a pipeline stage in the main processor with a pipeline stage in the coprocessor (or vice-versa). No such queue is described in Brown.

The Examiner tries to equate the claimed "token" with the commands associated with the spec-queue 75. These commands, generated by the I-BOX and E-BOX of the CPU 10, are used to ensure that ordering is preserved when accessing memory. How are those commands placed

in the synchronizing queue as tokens by the predetermined pipeline stage when processing a coprocessor instruction? Brown also does not teach the claimed partner pipeline stage then processing the coprocessor instruction upon receipt of the claimed token from the synchronizing queue so that the first and second pipelines are synchronized between the predetermined pipeline stage and the partner pipeline stage.

So Brown lacks coprocessor instructions from the same sequence of instructions as the main processor routed through both the pipelines of the main processor and the coprocessor. Further, Brown lacks a queue by which a synchronizing token is passed between the main processor and a coprocessor. In fact, Brown does not even consider the problem of synchronization between a main processor pipeline and a coprocessor pipeline in situations where a coprocessor instruction is routed through both pipelines.

The application is in condition for allowance. An early notice to that effect is earnestly solicited.

Respectfully submitted,

NIXON & VANDERHYE P.C.

By: 

John R. Lastova
Reg. No. 33,149

JRL:sd
901 North Glebe Road, 11th Floor
Arlington, VA 22203-1808
Telephone: (703) 816-4000
Facsimile: (703) 816-4100